

TITLE

COMPOSITIONAL SERVICE RESOURCE RESERVATION

5 CROSS REFERENCE TO RELATED APPLICATIONS

Not Applicable.

STATEMENT REGARDING FEDERALLY SPONSORED RESEARCH

10

Not Applicable.

BACKGROUND OF THE INVENTION

15

The present invention relates in general to allocating data processing resources within a digital network, and, more specifically, to using communications management tools to assemble compositional services employing multiple types of processing elements deployed within the digital network.

20

The operation of digital networks require significant investment in infrastructure; yet competitive pressures and the fast pace of technological change result in low profit margins. Naturally, network providers (such as Internet service providers and backbone providers) seek to minimize their equipment cost, optimize latency, and maximize scalability, robustness, and maintainability of their networks. Network architectures directed toward satisfying these goals often cause a tradeoff in network flexibility, such as the ability to reconfigure resources to provide different services.

25

For instance, it may be cost effective to employ a network architecture in which data processing resources are distributed in separate remote locations with several units of one type of processing component at one location (i.e., network node) and several units of another type of processing component at another node. Those

30

processing resources may be used to generate standalone services. For instance, a group of Class 4 call processing resources may be at one location, and a group of Service Control Point (SCP) processing resource may be at another. A compositional service is one that is generated from, or composed of, standalone services. Advanced
5 Intelligent Network (AIN) telephony services are an example of compositional services generated by combining the previously mentioned call processing and service control point processing services. Another such service is the delivery of audiovisual content (e.g., videoconferencing or streaming of movies) to an end user. Various component services will be necessary to provide the complete compositional service,
10 such as data storage services, media conversion services, authentication/authorization services, and billing services.

Physical processing components for performing the functions of a compositional service typically have been manually configured to compose the service. The capacity of the processing components dedicated to the composed service
15 have been based on the expected usage of the service and balanced against the demands generated by other components at expected usage levels. To guarantee a certain performance level (i.e., quality of service or QoS) at times of peak usage, a capacity level would have to be built into the services which is then unused at off-peak times (which is usually a majority of the time). To avoid excessive cost, a service
20 usually is provided with enough capacity to handle typical loads but less capacity than would be required to handle peaks loads without degradation. Thus, a component dedicated to one service may be unused or underutilized during a time that another service is congested because of a lack of capacity of that same type of component. In other words, a component of a type "X" included in a composed group of components
25 for a service of a type "A" may be unused because of low usage of service "A" while a nearby component of type "X" for another service "B" is overloaded and providing a poor QoS due to peak usage of service "B". Conventional service configuration methods have not allowed sharing or reconfiguration of such resources in a simple, cost-effective manner. For instance, heavy local call volume may overwhelm call

processing resources while SCP lookup processing resources – that are used to facilitate inter-exchange calls – are largely idle.

SUMMARY OF THE INVENTION

5

The present invention has the advantage of providing a simple, functional mechanism for implementing service resource allocation which merges transport, processing, and storage resource allocation. The invention advantageously allows transport management methods and tools to be leveraged to provide compositional service assembly and unified resource management.

10

Briefly, the invention provides a mechanism for handling processing (i.e., computing) capacity in a manner analogous to handling transport quality of service. The memory, speed of operations, and input/output rates of a resource component are viewed as an aggregate bandwidth of transport – in essence, establishing a sustainable data flow rate analogous to the sustainable data flow rate indicated by a bandwidth measure in transport. The distributed functional components of a compositional application/service are adapted to receive and process resource allocation requests. This allows the same network protocols and tools as are used for transport management to be used for composing (i.e., assembling) and allocating (i.e., assigning) computing resources with a guaranteed delivery rate.

15

20

A compositional service is advertised by some method, such as UDDI or Web Pop-up. The compositional service is composed of a plurality of functional components by a central resource aggregator. Each component may be managed as a functional instance, as a service component, as a component type, or as a member of a set of components connected by identified transport links.

25

In one aspect of the invention, a method is provided for managing computer processing resources connected within a network. A plurality of physical processing components are interconnected within the network for providing a plurality of virtual processing elements that are accessible by respective network traffic paths. A

resource aggregator represents a pool of said virtual processing elements, each virtual processing element having a capacity allocable according to a respective communication transfer rate. A reservation request is received for utilizing specified processing resources. The resource aggregator identifies at least one virtual processing element for providing capacity to satisfy the reservation request in response to the respective communication transfer rate. Use of a respective network traffic path is allocated to service the reservation request in response to the identified virtual processing element.

BRIEF DESCRIPTION OF THE DRAWINGS

Figure 1 is a block diagram showing the generation of virtual components within physical components interconnected according to the present invention.

Figure 2 is a block diagram of a resource aggregator interconnected with multiple component types in one preferred embodiment of the invention.

Figure 3 is a block diagram showing aggregation of a compositional service type according to the present invention.

Figure 4 is a flowchart of a preferred method for allocating virtual processing elements.

Figure 5 is a flowchart of one preferred method of assembling a compositional service.

DETAILED DESCRIPTION OF PREFERRED EMBODIMENTS

Referring to Figure 1, a user 10 is connected to a network 11 for accessing functional services composed of physical processing components 12 and/or 13. In order to manage the resources available from components 12 and 13, they divide their processing capacity into pluralities of virtual processing elements 14 and 15, respectively. Each virtual processing element is capable of performing a respective

processing operation and has a respective amount of capacity which is preferably measured in terms of an equivalent communication transfer rate based on the speed at which complete data processing transactions (i.e., from input of data to a corresponding output of data) are accomplished (preferably measured in bytes per second). For example, a transcoder component translates from G.711 encoded format to G.729 encoded format at 64kbps, so the component is registered in the resource pool as a 64kbps "G.711/G.729 transcoder link". Each of these virtual processing elements are separately allocable and made accessible by respective network traffic paths. The paths to virtual processing elements residing on the same physical device may be comprised of the same transport media but may be separately identified according to respective port assignments or packet labeling, for example.

Applications being provided by the physical components are preferably multi-threaded so that a resource can be shared by multiple users. Otherwise, the full capacity of a component may have to be allocated in its entirety to service a single request.

Network 11 may be comprised of an internet protocol (IP) network, an asynchronous transfer mode (ATM) network, or other type of digital network. It may further comprise an internetwork such as the Internet.

Upon being connected to network 11, physical components 12 and 13 advertise the virtual processing elements that it can produce (i.e., automatically send messages to announce their capacity to other devices within the network). As shown in Figure 2, the advertisements are directed to a resource aggregator 20 which may preferably be an application program executing on a specified server connected to network 11. The advertisements are provided according to predetermined component types, such as a storage type, a media conversion type, or an

authentication/authorization type, for example. There may be a separate aggregator process 21 for a component type "X" that represents the available type "X" virtual processing components within a type "X" pool 22. Another physical component 16 also advertising virtual processing elements 17 of type "X" are also included in pool 22 for allocation by aggregator 20. A separate aggregator process 23 for a component

type "Y" represents the available type "Y" virtual processing components within a type "Y" pool 24.

Aggregator 20 includes network links such as a link 25 for communicating with additional physical processing components advertising virtual processing elements, with system administrative resources, and with end users, for example. Based on the pooled resources, aggregator 20 advertises the virtual processing elements to potential users of the resources within the network. Reservation requests seeking use of the virtual processing elements may be received via link 25.

When a new physical component is added to the network, its capacity is automatically added to a corresponding pool as new virtual processing elements in the aggregator and becomes available to allocation in response to user reservation requests.

The virtual processing elements run independently of the type of hardware platforms they physically reside on and the type of network connectivity being employed. They are configured to operate according to predetermined specifications for each processing operation to be made available. By making each virtual processing element available as a virtual network element that can be individually addressed and reserved, integrated communication and management applications/tools can be used, especially where part of an IP network.

Figure 3 shows a preferred embodiment for assembling complex service functions based on combining multiple component types into composite resource sets. Thus, aggregator 20 includes a process 26 for aggregating a service type "A" using virtual processing elements from pool 22 of component type "X" and pool 24 of component type "Y" along with at least one transport link for joining the virtual processing elements together and/or with a requesting user. Each instance of aggregated service type "A" comprises a composite resource set which is made available in a pool 27 of service type "A". Thus, when a "service network" is started by a network provider, an aggregation component is started for each service type that is to be made available to users. Based on the pools of virtual processing elements of

different component types (i.e., processing operations), the service-type aggregation component establishes and advertises a service resource pool of available composite resource sets, each with a respective capacity that may be determined in advance or in response to specific reservation requests. Each composite resource set includes a specific set of virtual processing elements having predetermined interactions for integrating their processing operations into a particular service function, such as delivery of video content. At least upon allocation, each composite resource set preferably includes reserved transport links to be used between the assigned virtual processing elements and/or between the user and the virtual processing elements.

To preferred embodiment, compositional service resource reservation is performed by making the physical processing components receive and process communications bandwidth reservation protocols. When receiving a reservation request message, the resource aggregator first checks to see if the requested resource is available in sufficient quantity to fulfill the request. This is easily performed by considering the resource request in terms of the resources required to process data at the given bandwidth line rate. If sufficient resources are available, the aggregator allocates the necessary composite resource sets and labels the allocated resources with a request session identifier to show that they are then in use. If sufficient resources are not available, the aggregator replies to the request with a decline message. When the user no longer requires the resources, they are de-allocated and a communications reservation release message is sent to the physical processing components.

A first method of the present invention will be described in greater detail in connection with Figure 4. In step 30, physical resources are connected to the network. Each resource sets up virtual processing elements with respective throughput capacities in step 31. In step 32, the physical resources advertise their virtual processing elements and capacities to the aggregator. In step 34, the aggregator advertises the available resources within its pool(s).

In step 35, a reservation request is received by the aggregator for use of specified virtual processing element capacity. The aggregator checks in step 36 to

determine whether its pool contains available capacity to satisfy the request. A decision is made in step 37 whether capacity is available and can be allocated. If not, a decline message is sent in step 30 and a return is made to step 35 to wait for additional reservation requests.

5 If capacity is available, the necessary capacity is reserved by the aggregator in step 40. Reservations are preferably made for the virtual processing elements under control of a selected physical resource and for the transport links necessary for end-to-end service using the allocated virtual processing elements. In an IP network, a reservation protocol such as RSVP-TE may preferably be used. In an ATM network,
10 resource allocation may be performed using an ATM virtual path as is known in the art. After use of the chosen virtual processing element(s) and transport link(s) have been secured, the aggregator allocates the pooled resources in step 41 by setting in-use flags corresponding to the allocated resources. In step 42, the aggregator transmits connection information to the requesting user (i.e., it identifies the network traffic
15 paths corresponding to the virtual processing elements and transport links). The connection information may, for example, include a label-switched path in an IP network corresponding to the path taken by user data through the processing operations of the virtual processing elements and to the requesting user.

 When the requesting user no longer needs the allocated resources, they
20 preferably send a reservation release message which is received in step 43. In step 44, the aggregator de-allocates the corresponding virtual processing elements and then forwards reservation release messages to the physical components and to the transport links involved.

 A further method of the invention for forming compositional services is
25 shown in Figure 5. After the aggregator constructs respective pools for each component type (as shown by steps 30-33 in Figure 4), it advertises predetermined service types in step 50. In the embodiment shown in Figure 5, resources to fulfill a reservation requests are assembled on-demand. Alternatively, a predetermined amount

of resources could be pre-assembled according to the service type in order to reduce response time to a reservation request.

In step 51, the aggregator receives a reservation request for a particular service type. A check is made in step 52 to determine whether the component types and transport link capacity are available in order to satisfy the service request. If not, then a decline message is sent in step 53 and a return is made to step 51 to await further requests.

If sufficient resources are found in the corresponding pools, then the aggregator constructs a service resource set for the requested service type in step 54.

In step 55, the aggregator reserves capacity within the assembled composite resource set (e.g., by sending one or more RSVP-TE messages) in order to be able to guarantee a predetermined quality of service for the requested service function. The aggregator updates the pools in step 56 by flagging the allocated resources with a request session identifier. In step 57, the aggregator transmits composite resource set information such as a label-switched path to the requesting user, who benefits from the assembled service with a guaranteed delivery rate. Once the service is no longer required, it is released as shown by steps 43 and 44 in Figure 4.

The foregoing invention has provided a simple, functional mechanism for implementing service resource allocation which merges transport resource allocation with processing resource allocation and storage resource allocation. Transport management methods and tools such as Tivoli® NetView® from IBM Corporation and traffic management protocols such as RSVP-TE are leveraged to provide compositional service assembly and management while optimizing resource utilization.